

WHAT IS CLAIMED IS:

1: A data aggregation module comprising an aggregation engine and a multidimensional datastore, wherein the multidimensional datastore stores multidimensional data logically organized along N dimensions, and wherein the aggregation engine performs data aggregation operations on said multidimensional data by:

(a) performing a first stage of data aggregation operations along a first dimension of said N dimensions; and

(b) performing a second stage of aggregation operations for a given slice in the first dimension along another dimension of said N dimensions.

2. The data aggregation module of claim 1, wherein the second stage of aggregation operations involves recursive data aggregation operations for slices in said N dimensions.

3. The data aggregation module of claim 1, further comprising a data loading mechanism for loading data from a database, and a storage handler for storing in the multidimensional datastore the data loaded from the database and the aggregated data generated by the aggregation engine.

4. The data aggregation module of claim 1, further comprising an interface for receiving queries generated by a requestor, and control logic that, upon determining that the multi-dimensional datastore does not contain data required to service a given query, controls the aggregation engine to generate aggregated data required to service the given query and controls the aggregation module to return the aggregated data to the requestor.

5. The data aggregation module of claim 1, wherein said interface interfaces to an OLAP server (comprising OLAP analysis logic and presentation logic) and client machines operably coupled to the OLAP server to provide user-directed OLAP analysis, to thereby realize an OLAP system capable of performing data aggregation operations on the data, and storing and managing such data.

6. The data aggregation module of claim 3, wherein said client machines include a web-browser-based user interface that enables user access to the OLAP server.

7. The data aggregation module of claim 1, integral to a DBMS to thereby realize an improved DBMS capable of performing data aggregation operations on the data, and storing and managing such data.

8. The data aggregation module of claim 1, integral to a DBMS operably coupled to a plurality of client machines over a network, to thereby realize a data warehouse capable of performing data aggregation operations on the data, and storing and managing such data.

9. The data aggregation module of claim 1, wherein said interface implements a standard protocol for accessing data.

10. The data aggregation module of claim 9, wherein the standard protocol comprises one of OLDB, OLE-DB, ODBC, SQL, and JDBC.

11. The data aggregation module of claim 1, wherein the aggregation engine stores the resultant data of aggregation operations for the given slice as a record in a data file, wherein location of the record in the data file is stored in a directory.

12. The data aggregation module of claim 11, wherein the directory stores, for a given record, a start address and end address of the record and a physical address of the data file.

13. A method for data aggregation for use with a multidimensional datastore that stores multidimensional data logically organized along N dimensions, the method comprising the steps of:

(a) performing a first stage of data aggregation operations along a first dimension of said N dimensions; and

(b) performing a second stage of aggregation operations for a given slice in the first dimension along another dimension of said N dimensions; and

(c) storing resultant data in said multidimensional datastore.

14. The method of claim 13, wherein step (b) involves recursive data aggregation operations for slices in said N dimensions.

15. The method of claim 13, further comprising the step of loading data from a database, and storing the data loaded from the database in said multidimensional datastore.

16. The method of claim 13, further comprising the steps of:

receiving queries generated by a requestor; and

upon determining that the multi-dimensional datastore does not contain data required to service a given query, controlling the aggregation engine to generate aggregated data required to service the given query and returning the aggregated data to the requestor.

17. The method of claim 16, wherein the queries are received over an interface that implements a standard protocol for accessing data.

18. The method of claim 17, wherein the standard protocol comprises one of OLDB, OLE-DB, ODBC, SQL, and JDBC.

19. The method of claim 13, wherein the resultant data of aggregation operations for the given slice is stored as a record in a data file, wherein location of the record in the data file is stored in a directory.

20. The method of claim 19, wherein the directory stores, for a given record, a start address and end address of the record and a physical address of the data file.